

LSI DOCKET NO. 01-758

APPLICATION FOR LETTERS PATENT OF THE UNITED STATES

CERTIFICATE OF MAILING BY "EXPRESS MAIL"

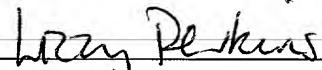
"EXPRESS MAIL" Mailing Label Number EL750736914US

Date of Deposit: December 6, 2001

I HEREBY CERTIFY THAT THIS CORRESPONDENCE, CONSISTING OF 17 PAGES OF SPECIFICATION AND 6 PAGES OF DRAWINGS, IS BEING DEPOSITED WITH THE UNITED STATES POSTAL SERVICE "EXPRESS MAIL POST OFFICE TO ADDRESSEE" SERVICE UNDER 37 CFR 1.10 ON THE DATE INDICATED ABOVE AND IS ADDRESSED TO: ASSISTANT COMMISSIONER OF PATENTS & TRADEMARKS, U.S. PATENT AND TRADEMARK OFFICE, P.O. BOX 2327, ARLINGTON, VA 22202.

BY:

Lizzy Perkins



SPECIFICATION

To all whom it may concern:

Be It Known, That We, **Charles E. Nichols**, a citizen of United States of America, residing at 4531 North Glendale, Wichita, Kansas 67220 and **Keith W. Holt**, a citizen of United States of America, residing at 1522 Krug Circle, Wichita, Kansas 67230, have invented certain new and useful improvements in "**Method and Apparatus to Manage Independent Memory Systems as a Shared Volume**", of which We declare the following to be a full, clear and exact description:

BACKGROUND OF THE INVENTION**1. Technical Field:**

The present invention is directed generally toward computer storage systems and, more particularly, toward a method and apparatus for managing independent storage controller memory systems as a single memory system for the purposes of allowing shared storage volume access.

2. Description of the Related Art:

Redundant Array of Independent Disks (RAID) is a disk subsystem that increases performance and provides fault tolerance. RAID is a set of two or more hard disks and specialized disk controllers that contain the RAID functionality. RAID can also be implemented via software only, but with less performance, especially when rebuilding data after a failure. RAID improves performance by disk striping, which interleaves bytes or groups of bytes across multiple drives, so more than one disk is reading and writing simultaneously. Fault tolerance is achieved by mirroring or parity. Mirroring involves duplication of the data on two drives. A failed drive can be hot swapped with a new one, and the RAID controller automatically rebuilds the lost data from the mirrored drive.

Dual, independent storage controllers are required to provide full data path redundancy to host computer systems. The controllers share access to the disk drives via their respective interface ports. The controllers present the data on the drives to one or more host systems as one or more logical volumes. However, simultaneous or interleaved access to data on a given volume from a plurality of controllers has associated cache coherency and data access latency problems. The coherency problems arise because each controller has an independent memory system for caching data from the volumes. Data access latency problems arise because the controllers must make their respective caches coherent when the two controllers interleave access to the data in the volumes.

One solution to the problems in the prior art is to not allow the controllers to simultaneously access the data. However, this approach restricts simultaneous data access to

hosts connected to a single controller. Another solution is to share a common data cache between a plurality of controllers. This approach is lacking because the common data cache is a single point of failure. Yet another solution is to establish an ownership model where controllers trade off the data access privileges. However, there are latencies associated with ownership transfer. These latencies are visible to the host computer systems.

Therefore, it would be advantageous to provide an improved method and apparatus for managing cache memory for a storage volume.

SUMMARY OF THE INVENTION

The present invention provides a switched architecture to allow controllers to manage physically independent memory systems as a single, large memory system. The switched 5 architecture includes a path between switches of controllers for inter-controller access to memory systems and input/output interfaces in a redundant controller environment. Controller memory systems are physically independent of each other; however, they are logically managed as a single, large memory pool. Cache coherency is concurrently maintained by both controllers through a shared locking mechanism. Volume Logical Block Address extents or individual cache blocks can be locked for either shared or exclusive access by either controller. There is no strict ownership model to determine data access. Access is managed by the controller that receives the access request. When a controller is removed or fails, a surviving controller may take appropriate action to invalidate all cache data that physically resides in the failed or missing controller's memory systems. Cached write data will be mirrored between redundant controllers to prevent a single point of failure with respect to unwritten cached write data.

10
15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself however, as well as a preferred mode of use, further objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Figure 1 is a block diagram of a switched controller architecture in accordance with a preferred embodiment of the present invention;

Figures 2A and 2B are block diagrams illustrating example shared cache read sequences in accordance with a preferred embodiment of the present invention;

Figures 3A and 3B are block diagrams illustrating example shared cache write sequences in accordance with a preferred embodiment of the present invention;

Figure 4 is a flowchart illustrating the processing of a read request in accordance with a preferred embodiment of the present invention; and

Figure 5, a flowchart illustrating the processing of a write request is shown in accordance with a preferred embodiment of the present invention.

DETAILED DESCRIPTION

The description of the preferred embodiment of the present invention has been presented for purposes of illustration and description, but is not limited to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art. The embodiment was chosen and described in order to best explain the principles of the invention the practical application to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

With reference now to the figures and in particular with reference to **Figure 1**, a block diagram of a switched controller architecture is depicted in accordance with a preferred embodiment of the present invention. The architecture includes a first controller **100** and a second controller **150** to provide full path redundancy to host computer systems. Controller **100** includes host channel adapters (CA) **102, 104** and drive channel adapters **106, 108**. The host channel adapters are the physical connections between the internal bus and the host interface. The internal bus may be, for example, an Infiniband bus. While the example shown in **Figure 1** is an Infiniband architecture, controllers **100, 150** may be any other switched architecture. The drive channel adapters are the physical connections between the internal bus and the drive interface.

Controller **100** also includes central processor unit (CPU) **110**. The CPU may have an associated random access memory (RAM) **112** as a working memory. Further, controller **100** includes remote memory controllers (RMC) **122, 124**. A RMC is the control hardware for managing the connection to a memory. RMC **122** manages the connection to RAM **126** and RMC **124** manages the connection to RAM **128**.

Host channel adapters **102, 104**, drive channel adapters **106, 108**, CPU **110**, and remote memory controllers **122, 124** are connected using switch **130**. The switch is a semi-intelligent hardware component with multiple ports. A request received on any port can be directly routed to any other port on the switch. In the example of an Infiniband controller, switch **130** is an Infiniband switch.

Controller 150 includes host channel adapters 152, 154 and drive channel adapters 156, 158. Controller 150 also includes CPU 160. The CPU may have an associated random access memory (RAM) 162 as a working memory. Further, controller 150 includes remote memory controllers 172, 174. RMC 172 manages the connection to RAM 176 and RMC 174 manages the connection to RAM 178. Host channel adapters 152, 154, drive channel adapters 156, 158, CPU 160, and remote memory controllers 172, 174 are connected using switch 180.

In accordance with a preferred embodiment of the present invention, the switched architecture includes path 190 between switch 130 and switch 180. Path 190 is a switch-to-switch path that allows for inter-controller access to memory systems and input/output (I/O) interfaces in a redundant controller environment. For example, when a request is received on host CA 102, CPU 110 may access a device via drive CA 156 through path 190. As a further example, when a request is received on host CA 154, CPU 160 may access RAM 128 via RMC 124 through path 190.

Switch-to-switch path 190 may be provided through edge connectors. Alternatively, path 190 may be provided through a wired connection between controller cards. Other techniques for providing path 190 may also be used within the scope of the present invention.

Those of ordinary skill in the art will appreciate that the hardware depicted in **Figure 1** may vary. For example, each controller may include more or fewer host channel adapters. The controllers may also include more or fewer drive channel adapters depending on the implementation. While the example depicted in **Figure 1** shows two memory controllers 122, 124 and two random access memories 126, 128, more or fewer memories and associated controllers may be used. In addition, a controller may include a plurality of central processing units. The depicted example is not meant to imply architectural limitations with respect to the present invention.

The controller memory systems are physically independent of each other. However, according to a preferred embodiment of the present invention they are logically managed as a single, large memory pool. Cache coherency is concurrently maintained by both controllers through a shared locking mechanism. Volume Logical Block Address (LBA) extents or individual cache blocks can be locked for either shared or exclusive access by either controller.

There is no strict ownership model to determine data access. Access is managed by the controller that receives the access request. When a controller is removed or fails, a surviving controller may take appropriate action to invalidate all cache data that physically resides in the failed or missing controller's memory systems. Cached write data may be mirrored between redundant controllers to prevent a single point of failure with respect to unwritten cached write data.

With reference now to **Figures 2A** and **2B**, block diagrams illustrating example shared cache read sequences are shown in accordance with a preferred embodiment of the present invention. Particularly, with respect to **Figure 2A**, a read request is processed according to the following steps:

1. A read request is received by Controller A.
2. Controller A allocates memory buffers for the read data. Because the logical cache memory pool resides on both controllers, the memory buffer could be allocated from either controller's physical memory pool. In the example shown in **Figure 2A**, the buffer for the read request received by Controller A happens to be on Controller B. By definition, however, this buffer could be allocated on either controller. It is during this allocation phase that cache coherency must be maintained between the controllers. The LBA extent for this read is marked as locked, such that other reads to the same LBA on either controller is forced to wait for the disk read to complete for the read received by Controller A. Once the memory buffer is allocated, Controller A maps the request to the appropriate disk drives and initiates reads (data transfers) via the appropriate drive CA from the disk drives. The reads do not necessarily have to occur through the drive CA on the controller that received the original read request.
3. The drive CA begins to transfer the data to the appropriate memory pool. This step facilitates future cache read hits for this data. Because of the concurrent cache coherency management inherent in this approach, subsequent reads of the same LBA to either controller would discover the data in the logical cache pool.
4. Data is transferred to the host CA on Controller A that received the request.

5. Controller A directs command status to be returned through the originating CA on that controller.

Turning now to **Figure 2B**, a read request for which the data exists in cache is processed according to the following steps:

- 5 1. A read request is received by Controller A.
2. Controller A discovers the LBA in the logical cache pool and initiates reads from the memory on Controller B.
3. The RMC on Controller B reads the data from memory.
4. Data is transferred to the host CA on Controller A that received the request.
- 10 5. Controller A directs command status to be returned through the originating CA on the controller.

15 With reference now to **Figures 3A** and **3B**, block diagrams illustrating example shared cache write sequences are shown in accordance with a preferred embodiment of the present invention. Particularly, with respect to **Figure 3A**, a write request is processed with write-back caching according to the following steps:

- 20 1. A write request is received by Controller A. Controller A allocates memory buffers for the request. Two buffers are allocated, one on Controller A and another on Controller B. These two buffers serve as mirrors of each other. In order to maintain cache coherency, the LBA extent is locked to prevent other access to this data by requests received by either controller.
2. The data transfer is initiated by Controller A. The originating host CA begins to transfer data to the primary data buffer via the appropriate RMC. Although the example in **Figure 3** shows that the primary data buffer resides on Controller A, the primary data buffer may reside on either controller. However, it is required that the mirror buffer reside on the controller that does not contain the primary data buffer.
- 25 3. Data is transferred to the appropriate RMC and data buffer on the alternate controller.
4. Controller A directs command status to be returned through the originating CA.

Turning now to **Figure 3B**, a write request is processed with write-through caching according to the following steps (note that write cache mirroring may not occur during write-through requests):

1. A write request is received by Controller A. Controller A allocates memory buffers for the request. (Optional) Two buffers are allocated, one on Controller A and another on Controller B. These two buffers serve as mirrors of each other. In order to maintain cache coherency, the LBA extent is locked to prevent other access to this data by requests received by either controller.
2. The data transfer is initiated by Controller A. The originating host CA begins to transfer data to the primary data buffer via the appropriate RMC.
3. (Optional) Data is transferred to the appropriate RMC and data buffer on the alternate controller.
4. Controller A directs write completion to the disk drives.
5. Controller A directs command status to be returned through the originating CA.

Next, with reference to **Figure 4**, a flowchart illustrating the processing of a read request is shown in accordance with a preferred embodiment of the present invention. The process begins when a read request is received and a determination is made as to whether the data block is cached (step 402). If the data block is not cached, the process allocates memory buffers for read data (step 404), accesses the data block on the drive (step 406), and stores the data block in the memory pool (step 408). The data block may be accessed through a drive CA on any controller connected through the switch path. Furthermore, the memory buffers may be allocated and stored on a memory system of any controller connected through the switch path.

Then, the process transfers data to the host CA (step 410), returns command status (step 412), and ends. If the data block is cached in step 402, the process access the data block in the memory pool (step 414) and proceeds to step 410 to transfer the data to the host CA, return command status (step 412), and end.

Turning now to **Figure 5**, a flowchart illustrating the processing of a write request is shown in accordance with a preferred embodiment of the present invention. The process begins when a write request is received and allocates memory buffers for write data (step 504) and

transfers data to the primary data buffer (step 506). Thereafter, the process transfers data to the mirror data buffer (step 508). The primary buffer need not reside on the controller that receives the write request. However, the mirror buffer must reside on a controller, which does not contain the primary buffer to avoid a single point of failure. Next, the process returns command status 5 (step 510) and ends. Alternatively, if the volume is configured for write-through caching, the controller directs write completion to the disk drives before returning status in step 510.

Thus, the present invention solves the disadvantages of the prior art by utilizing the switched architecture of the controllers to treat physically independent memory systems as a single, large logical memory system. The switched architecture facilitates direct data transfers to components that are not on board with respect to a single controller. From a host perspective, this approach eliminates a strict ownership model within a redundant controller storage environment. A host can access data from either storage controller without being exposed to the ownership change latency associated with moving ownership between controllers. Because there are no preferred access paths, I/O performance to a given volume is nearly identical on both controllers, thus eliminating latency involved in directing access from a non-preferred controller to a preferred controller. The present invention also provides a shared cache system without excess latency. The shared cache volume also is not a single point of failure, because it allows mirroring between independent memory systems.

10
15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95